arXiv:1410.3764v2 [cs.DS] 3 Dec 2015

A LAZY APPROACH TO ON-LINE BIPARTITE MATCHING

JAKUB KOZIK AND GRZEGORZ MATECKI

Abstract. We present a new approach, called a lazy matching, to the problem of on-line matching on bipartite graphs. Imagine that one side of a graph is given and the vertices of the other side are arriving on-line. Originally, incoming vertex is either irrevocably matched to an another element or stays forever unmatched. A lazy algorithm is allowed to match a new vertex to a group of elements (possibly empty) and afterwords, forced against next vertices, may give up parts of the group. The restriction is that all the time each element is in at most one group. We present an optimal lazy algorithm (deterministic) and prove that its competitive ratio equals $1 - \pi/\cosh(\frac{\sqrt{3}}{2}\pi) \approx 0.588$. The lazy approach allows us to break the barrier of $1/2$, which is the best competitive ratio that can be guaranteed by any deterministic algorithm in the classical on-line matching.

1. Introduction

Many problems of task-server assignment can be modeled as finding a matching in a bipartite graph $G = (U, D, E)$. Vertices of one part (set D) correspond to servers, and the vertices of the other (set U) – to tasks. An edge between a task and a server indicates that the server is capable of performing the task. In a simple setting, when one server can realize at most one task, the problem of maximization of the number of realized tasks reduces to finding a maximum matching. In real-life applications it is very common that not all tasks are known a priori and some decisions about assignments have to be taken with no knowledge about future tasks. A simple model for this situation is on-line bipartite matching. In this setting servers are known from the beginning and tasks are revealed one by one. The decision about assignment of each task has to be made just after its arrival and cannot be changed in the future. Suppose that there are n servers, n tasks are going to be revealed, and capabilities of servers are such that it is possible to realize all the tasks. (i.e. there exists a perfect matching in the tasks-servers graph). It is an easy exercise to show that even with these restrictions it is possible to present tasks in such a way, that the constructed assignment is at most $\lceil n/2 \rceil$. On the other hand, any greedy assignment strategy guarantees that at least half of the tasks will be assigned.

In their classical contribution Karp, Vazirani, and Vazirani [14] take other approach in which the graph to be presented is fixed before the first task is presented. In particular the presented graph does not depend on the decisions of the assigning algorithm. It does

Theoretical Computer Science, Faculty of Mathematics and Computer Science, Jagiellonian University in Kraków, Poland

E-mail addresses: Jakub.Kozik@tcs.uj.edu.pl, Grzegorz.Matecki@tcs.uj.edu.pl.

Key words and phrases. on-line, bipartite matching, adaptive algorithm.

Research of J. Kozik was supported by Polish National Science Center UMO-2011/03/D/ST6/01370.

Research of G. Matecki was supported by Polish National Science Center 2011/03/B/ST6/01367.

not make any difference for the worst-case analysis of the algorithm, but the approach provides a framework for analyzing randomized ones. The authors presented a randomized algorithm which on the average constructs a matching of size at least $(1 - 1/e)n$. They also argued that the result is asymptotically best possible for any randomized algorithm (the original paper [14] contained a mistake, which has been corrected in [11], see also a simplified exposition [4]). The approach of [14] has been applied to many variants of the original problem, with various practical applications (switch routing problem [1, 2], on-line auctions [16], Adwords problem [6, 11, 18] etc.) Recently a lot of interest is put into a problem of on-line stochastic matching [3, 9, 13, 16, 17] where a competitive ratio can be greater than $1 - \frac{1}{e}$. A different approach (called b-matching) is presented in [12] where authors allow a server to realize up to b tasks at the same time. They showed an optimal deterministic algorithm with competitive ratio $1 - \frac{1}{(1+\frac{1}{b})^b}$ (which tends to $1 - \frac{1}{e}$ with $b \to \infty$).

It is not hard to imagine a situation that the cost of a running server is roughly the same as the cost of an idle server. It might be profitable to start to realize some task on many servers simultaneously. In this work, we allow our algorithm to assign more than one server to an incoming task, and later on gradually giving up its computations and switch some servers to yet another tasks. Clearly, it is sufficient that at least one server completes the computations of a task. This enables to postpone the decision about which server is going to accomplish a given task. With this settings we present an algorithm matching at least $(1 - \pi/\cosh(\frac{\sqrt{3}}{2}\pi))n + o(n) \approx 0.588n + o(n)$ vertices where n is the maximum size of a matching in the whole presented bipartite graph.

The lazy approach was first introduced by Felsner in [10] as an adaptive generalization of the on-line chain partitioning problem. His adaptive on-line modification is in our terms a lazy approach to chain partitioning of posets (see also [15]). It is still open and seems challenging to verify if adaptive (lazy) approach to chain partitioning allows more efficient on-line algorithms.

1.1. **Related Work.** The problem of b-matching seems similar to lazy matching only the roles of servers and tasks are switched. However, in the lazy approach an algorithm always ends up with matching of type one to one, while a b-matching algorithm can assign one server to many tasks.

Another similar approach was proposed by Feldman et al. [8] as *free disposal*. They consider the weighted matching problem where each incoming vertex $u \in U$ may be assigned to one of its neighbors or left alone. Each vertex $d \in D$ accepts at most $n(d)$ vertices from U with highest-weighted edges. Here roles of servers and tasks are switched. All tasks are given at once and servers are incoming on-line. Each server has to be assigned to at most one task. In the end, each task chooses at most $n(d)$ servers from all the servers assigned to it – the ones with the highest-weighted edge. The main difference from the lazy approach is that once a connection between a server and a task is established it cannot be changed till the very end. There is no such restriction in the lazy approach – a server may drop its task and take a new one during the on-line process.

The idea of dropping an edge from already constructed matching is investigated in the *preemptive model*. Here, edges with weights are incoming on-line and algorithm is

allowed to remove previously accepted edges in order to add a new one. A collection of results on preemptive matching can be found in [5, 7].

1.2. **Problem definition.** For a positive integer α, the *α-lazy matching game* is played in rounds between Scheduler and Builder. They play on a set of vertices D, known in advance. There is finite number of rounds and in each round:

(1) Builder presents a vertex u and reveals its neighbors $N(u) \subseteq D$.
(2) Scheduler assigns to u a set $m(u) \subseteq N(u)$ of size at most α and updates $m(x) = m(x) \setminus m(u)$ for every vertex x presented before.

Let U be the set of all vertices presented by Builder in the game and let $G = (U, D, E)$ be an underlying bipartite graph with $(u, d) \in E$ when $d \in N(u)$, for all $u \in U, d \in D$. Now, the *size* n of the game is the maximum size of a matching in G. If either $\alpha \geqslant |D|$ for all n or more generally $\alpha \overset{n\to\infty}{\longrightarrow} \infty$ then the game is called *∞-lazy matching* or simply *lazy matching*. We refer to this by writing $\alpha \to \infty$.

The goal of Scheduler is to maximize the number of nonempty sets $m(u)$ over all vertices $u \in U$. Intuitively, such vertex u is successfully matched with arbitrary chosen $d \in m(u)$. We refer to this number as *the size of the matching constructed by Scheduler*. The goal of Builder is just the opposite, he disturbs Scheduler as much as he can.

The interpretation of α-lazy game into servers-tasks assignment is clear: D is the set of servers, U is the set of incoming tasks and algorithm assigns servers in $m(u)$ to an incoming u (possibly canceling previous computations on servers in $m(u)$). The quality of Scheduler is measured by the number (or the fraction) of tasks which are being realized at the end of the game. Note that for $\alpha = 1$ the game reduces to the classical on-line bipartite matching.

Let $\mathcal{A}$ be an algorithm which assigns incoming tasks. We denote by $\mathrm{val}_{\mathcal{A}}(n)$ the worst case value of the matching constructed by $\mathcal{A}$ in all possible games of size n. The value of the α-lazy matching problem $\mathrm{val}_\alpha(n)$ is the maximum value of $\mathrm{val}_{\mathcal{A}}(n)$ among all α-assigning algorithms $\mathcal{A}$. Since no algorithm produces matching larger that n we additionally use a *competitive ratio* defined as $\liminf_{n\to\infty} \mathrm{val}_\alpha(n)/n$.

1.3. **Main results.** To solve the problem of α-lazy on-line matching we consider a deterministic algorithm, called α-BALANCED. The algorithm is described in Section 3. It behaves greedy, i.e., no task is rejected if there is a possibility to run it, and tries to locally balance the sizes of all $m(u)$. We prove α-BALANCED algorithm is the best possible one.

Theorem 1.1. *α-BALANCED is an optimal strategy for Scheduler in the α-lazy matching game.*

The next two sections are to prove the theorem. The proof is split into two parts. The following schema of system of inequalities is crucial for both arguments:

$$\begin{cases} (1+\alpha)\mathbf{x_0} \leqslant n, \\ (\mathbf{x_0} + \ldots + \mathbf{x_i})(1 + \mathbf{x_i}) \leqslant n - i, \quad i = 1, \ldots, k, \\ \mathbf{x_1} \geqslant \mathbf{x_2} \geqslant \ldots \geqslant \mathbf{x_k} \geqslant 0, \\ \mathbf{x_0} + \ldots + \mathbf{x_k} \geqslant 0. \end{cases} \tag{1}$$

We are going to work with n and α fixed. Then, the schema is parametrized by a positive integer k. We say that a pair (k, x) satisfy system (1) if $x = (x_0, x_1, \ldots, x_k)$ is an integer vector satisfying instance of the schema for this particular k.

In Section 2, we prove (Proposition 2.2) that for every solution (k, x) of (1), every α-lazy algorithm $\mathcal{A}$ can be cheated by some Builder's strategy in a game of size n so that $\mathcal{A}$ matches at most $n - (x_0 + \ldots + x_k)$ vertices. On the other hand, in Section 3, we show (Proposition 3.1) that when α-BALANCED constructs a matching of size k in a game of size n, then $k = n - (x_0 + \ldots + x_k)$ for some (k, x) satisfying (1). These two facts ensure α-BALANCED is the optimal strategy for Scheduler.

In order to determine the competitive ratio of α-BALANCED we need to maximize the sum $x_0 + \ldots + x_k$ over all feasible solutions (k, x) of (1). In Section 4 we present the linear programming formulation of (1) which, with the use of the Complementary Slackness Theorem, can be easily solved. Finally, we prove

Theorem 1.2. *The competitive ratio of α-lazy on-line matching problem on bipartite graphs (and the competitive ratio of α-BALANCED algorithm) equals $1 - \frac{\alpha}{1+\alpha} \prod_{i=1}^{\alpha-1} \frac{i+i^2}{1+i+i^2}$. For $\alpha \to \infty$ it converges to $1 - \pi/\cosh \frac{\sqrt{3}\pi}{2} \approx 0.588$.*

The ratio converges fast and for $\alpha \in \{2, 3\}$ it is almost optimal, i.e., $5/9 \approx 0.556$ and $4/7 \approx 0.571$, respectively.

2. Worst case scenario for a lazy algorithm

Inequalities (1) allow x_0 to be negative. We start with an observation that in order to maximize the sum $x_0 + \ldots + x_k$ it suffices to consider solutions of (1) with $x_0 = \left\lfloor \frac{n}{1+\alpha} \right\rfloor$.

Proposition 2.1. *For any pair (k, x) satisfying (1) there exists a pair (k', x') satisfying (1) such that $x'_0 = \left\lfloor \frac{n}{1+\alpha} \right\rfloor$ and $x'_0 + \ldots + x'_{k'} \geqslant x_0 + \ldots + x_k$.*

Proof. The case when $x_1 = 0$ is easy: just put $x'_0 = \left\lfloor \frac{n}{1+\alpha} \right\rfloor$, $k' = 1$ and $x'_1 = 0$. We assume that $x_1 > 0$. The claim is proved by induction on $\left\lfloor \frac{n}{1+\alpha} \right\rfloor - x_0$. The base, when $x_0 = \left\lfloor \frac{n}{1+\alpha} \right\rfloor$, is obvious. For the induction step, let us assume that $x_0 < \left\lfloor \frac{n}{1+\alpha} \right\rfloor$. Let j be the greatest index for which $x_j = x_1$. Consider the following sequence $(x'_0, \ldots, x'_k)$: $x'_0 = x_0 + 1$, $x'_j = x_j - 1$ and $x'_i = x_i$ for all $i \notin \{0, j\}$. For $i \geqslant j$ we have $x'_0 + \ldots + x'_i = x_0 + \ldots + x_i$. Therefore

$$(x'_0 + \ldots + x'_i)(1 + x'_i) \leqslant (x_0 + \ldots + x_i)(1 + x_i) \leqslant n - i.$$

For $0 < i < j$ we get

$$\begin{aligned}(x'_0 + \ldots + x'_i)(1 + x'_i) &= (x_0 + \ldots + x_i + 1)(1 + x_j) \leqslant \\ &\leqslant (x_0 + \ldots + x_{j-1} + x_j)(1 + x_j) \leqslant \\ &\leqslant n - j \leqslant n - i.\end{aligned}$$

Since $x'_0 \leqslant \left\lfloor \frac{n}{1+\alpha} \right\rfloor$ sequence $(k, (x'_0, \ldots, x'_k))$ satisfies (1). Finally, $\left\lfloor \frac{n}{1+\alpha} \right\rfloor - x'_0 < \left\lfloor \frac{n}{1+\alpha} \right\rfloor - x_0$, so by the induction hypothesis there exists a solution of (1) satisfying the claim. □

Proposition 2.2. *For any pair (k, x) satisfying (1) there exists a strategy for Builder in the α-lazy matching game of size n such that any Scheduler constructs a matching of size at most $n - (x_0 + \ldots + x_k)$.*

Proof. By Proposition 2.1 it is enough to consider pairs with $x_0 \geqslant 0$. Without loss of generality we assume that $x_k > 0$ and describe a strategy for Builder, that does not allow Scheduler to construct a matching larger than $n - (x_0 + \ldots + x_k)$. During the game Builder presents a bipartite graph $G = (U, D, E)$ with $|U| = |D| = n$ and maintains an auxiliary structure: a partition of $U = U_0 \cup U_1 \cup \ldots \cup U_k \cup R$ and a partition of $D = D_0 \cup D_1 \cup \ldots \cup D_k \cup S$ such that

$$|U_0| = |D_0| = x_0,$$
$$|U_i| = |D_i| = 1 + x_i, \qquad \text{for } i = 1, \ldots, k,$$
$$N(u_i) = D - (D_0 \cup \ldots \cup D_{i-1}), \qquad \text{for each } u_i \in U_i, \tag{2}$$
$$N(r) = S, \qquad \text{for each } r \in R. \tag{3}$$

Observe that (1) guarantees that $x_0 + \ldots + x_k \leqslant n - k$ and thus $\sum_{i=0}^{k} |U_i| = \sum_{i=0}^{k} |D_i| \leqslant x_0 + \sum_{i=1}^{k}(1 + x_i) \leqslant n$. Therefore $|R| = |S| \geqslant 0$. It is straightforward that any bipartite graph which can be partitioned in such way contains a perfect matching.

The strategy of Builder is divided into $k + 2$ phases enumerated from 0 to $k + 1$. Figure 1 depicts the evaluation of the strategy described below for $\alpha = 2$, $n = 18$ and $k = 2$. In the beginning of the i-th phase ($0 \leqslant i \leqslant k$) sets U_j and D_j, for $j < i$, are already fixed. Next, during the i-th phase, Builder presents $1 + x_i$ vertices, or x_0 vertices when $i = 0$, which form set U_i with neighborhoods defined by (2). Most important, Builder chooses in the special way the set $D_i \subseteq D - (D_0 \cup \ldots D_{i-1})$ of size $1 + x_i$ when $i > 0$ and of size x_0 when $i = 0$. This will conclude phase i. At the very end, within phase $k+1$, Builder presets a set R of size $n-k-(x_0+\ldots+x_k)$ with vertices neighbouring with all vertices in $S = D - \bigcup_{i=0}^{k} D_i$.

It remains to define Builder's choice of D_i. For that, after each phase i ($0 \leqslant i \leqslant k$) Builder maintains the following:

($\star$) there are i distinct vertices $y_1, \ldots, y_i \in U_0 \cup \ldots \cup U_i$ such that $m(y) \cap \bigcup_{j=0}^{i} D_j = \emptyset$ for all $y \in (U_0 \cup \ldots \cup U_i) \setminus \{y_1, \ldots, y_i\}$.

Observe that, for a fixed $X \subseteq D$ and a fixed $y \in U$, once the condition $X \cap m(y) = \emptyset$ is satisfied, it will stay so to the end of the game as $m(y)$ may only shrink later on in the game. We prove by induction that choosing such D_i is possible in every phase by successively finding correct y_i's.

For $i = 0$, Builder chooses[1] any $D_0 \subseteq D - m(U_0)$ of size x_0 and it is possible as $|D - m(U_0)| \geqslant n - \alpha x_0 \geqslant x_0$ (because $|m(y)| \leqslant \alpha$ for all $y \in U_0$, and by (1)). This way $D_0 \cap m(y) = \emptyset$ for all $y \in U_0$ as required.

For $1 \leqslant i \leqslant k$, we assume that ($\star$) holds with vertices $Y = \{y_1, \ldots, y_{i-1}\}$ after $(i-1)$-th phase. Let $U' = U_0 \cup \ldots \cup U_i$ and $D' = D - (D_0 \cup \ldots \cup D_{i-1})$. First, note that $m(U' - Y) \subseteq D'$ by the ($\star$)-property after $(i-1)$-th phase. We split vertices in D' into two blocks: $D' = m(U' - Y) \cup X$, where X is simply the set of all unmatched vertices

[1]Let $m(X) = \bigcup_{x \in X} m(x)$ for a set X.

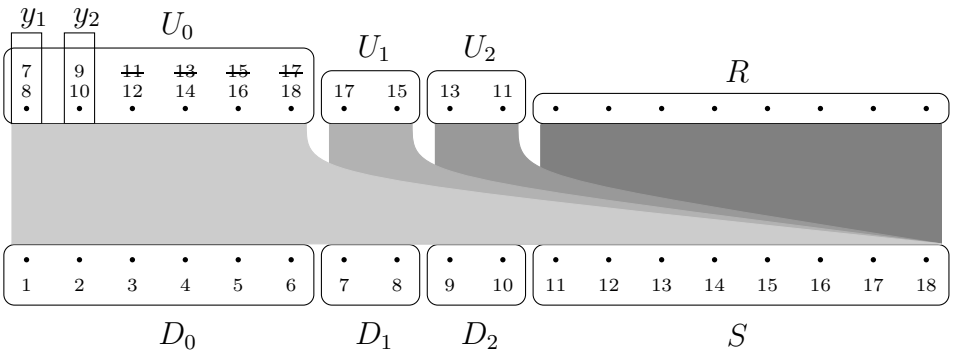

FIGURE 1. An example of the construction for solution $x = (6, 1, 1)$ of (1) with $n = 18$ and $\alpha = 2$. Recall that $N(U_0) = D$, $N(U_1) = D - D_0$, $N(U_2) = D_2 \cup S$ and $N(R) = S$. The numbers under vertices in the bottom part represent vertices' names, while the numbers in the upper part correspond to sets $m(u)$ (e.g. $m(y_1) = \{7, 8\}$).

in D'. Next, we give a lower bound on the average size of $m(u)$ for $u \in U' - Y$:

$$\begin{aligned}\frac{|m(U'-Y)|}{|U'-Y|} = \frac{|D'|-|X|}{|U'|-|Y|} &= \frac{n-i-|X|-(x_0+\ldots+x_{i-1}-1)}{x_0+\ldots+x_i+1} \\ &\overset{(1)}{\geqslant} \frac{(x_0+\ldots+x_i)(1+x_i)-|X|-(x_0+\ldots+x_{i-1}-1)}{x_0+\ldots+x_i+1} \\ &= x_i + \frac{1-|X|}{x_0+\ldots+x_i+1} > x_i - |X|\,.\end{aligned}$$

Clearly, there must be a vertex $y_i \in U' - Y$ with $|m(y_i)| \geqslant \frac{|m(U'-Y)|}{|U'-Y|} > x_i - |X|$. Since all values are integers we have $|X| + |m(y_i)| \geqslant 1 + x_i$. Builder picks D_i to be any subset of $X \cup m(y_i)$ of size $1 + x_i$, to keep the property $(\star)$ satisfied after the i-th phase.

By the condition $(\star)$ after the k-th phase there is $Y = \{y_1, \ldots, y_k\}$ such that $D - S$ and $m(y)$ are disjoint for all $y \in U - Y$. Therefore, whenever $m(u) \neq \emptyset$, for $u \in U$, then either $u \in Y$ or $m(u) \subseteq S$. It means that the number of such u's is at most $|Y| + |S| = k + n - (|D_0| + \ldots + |D_k|) = n - (x_0 + \ldots + x_k)$. Consequently, the size of the matching produced by Scheduler is at most $n - (x_0 + \ldots + x_k)$. □

3. The best matching algorithm

At any moment during a lazy matching game, we say that $d \in D$ is *available* for $u \in U$ if $d \in N(u)$ and $m(x) \neq \{d\}$ for any x presented earlier. Also, d is *strongly available* for u if it is available for u and d does not belong to any $m(x)$. Vertex $e \in U$ is *ready* for u if $m(e)$ contains an element which is available for u.

We present an algorithm for Scheduler called α-BALANCED. Suppose that vertex u has just been presented and let U be the set of vertices presented so far (including u). Each set $m(x)$ for $x \in U - \{u\}$ is already known and the algorithm has to construct set $m(u)$. The construction is described below – $m(u)$ is increased, one element at a

time, until certain condition is satisfied. During the process some other sets $m(x)$ may be decreased.

Algorithm 1: $\alpha -$ BALANCED(u)

1 **let** $m(u) := \emptyset$
2 pick up at most α strongly available elements for u and put it into $m(u)$
3 **while** *there exists a vertex* $e \in U$ *that is ready for* u *and satisfies*

$$|m(u)| + 2 \leqslant |m(e)|$$

do
4 from the set of all such vertices pick e with maximal size of $m(e)$
5 move one vertex available for u from $m(e)$ to $m(u)$

The condition in line 2 guarantees that the size of $m(u)$ will be at most α. Note that α-BALANCED never leaves $m(u)$ empty if there exists an available element for u, so in this respect α-BALANCED can be considered as greedy. For $\alpha = 1$ the above algorithm is just a simple greedy construction of a bipartite matching.

The following proposition describes the performance of α-BALANCED.

Proposition 3.1. *The size k of matching produced by α-BALANCED in a lazy matching game of size n equals $n-(x_0+x_1+\ldots+x_k)$ for some pair $(k,(x_0,\ldots,x_k))$ satisfying* (1).

Proof. Consider an instance of the lazy matching game of size n in which Builder produced graph $G = (U, D, E)$, and algorithm α-BALANCED constructed matching $m : U \to \mathcal{P}(D)$. Suppose that N rounds have been played in the game. *Presenting time* of element $u \in U$ is the index of the round in which u has been presented.

We denote by $m^t : U \to \mathcal{P}(D)$, the (partial) matching constructed up to round t. In particular, for $u \in U$ that is presented in round t_0, we have $m^t(u) = \emptyset$ for $t < t_0$ and then $(m^{t_0}(u), m^{t_0+1}(u), \ldots, m^N(u))$ is a weakly decreasing sequence of sets with $m^N(u) = m(u)$.

Let X be the set of all vertices in D such that $m(u) \cap X = \emptyset$ for each $u \in U$. The size of the matching produced be α-BALANCED is equal to the size of $\mathbf{Y} = \{u \in U : m(u) \neq \emptyset\}$. For the proof of the proposition we need the following claims.

Claim 3.2. Suppose that $x \neq y$, $m^{t_1}(x) \cap m^{t_2}(y) \neq \emptyset$ and $t_1 < t_2$, then $|m^{t_1+1}(x)| \geqslant |m^{t_2}(y)|$.

Proof. It is sufficient to verify the claim for $t_2 = t_1 + 1$. It means that during round t_2 algorithm α-BALANCED removed one element from $m(x)$ and inserted it into $m(y)$. Let $d \in m^{t_1}(x) \cap m^{t_2}(y)$ be the last such element (see Figure 2). This happens only when the condition from the line 3 of the algorithm is satisfied and x is a vertex with maximum size of assigned set among vertices ready for y. Let s be the size of the set assigned to x at that moment (in terms of listing 1 it is $|m(x)|$). Clearly $|m^{t_1}(x)| \geqslant s$, since $|m^{t_1}(x)|$ is the size of the set assigned to x in the beginning of round t_2, and that set can only get smaller during the round. Also $s - 1 = |m^{t_1+1}(x)|$ since element d was the last one removed from $m(x)$.

The condition from the line 3 of α-BALANCED guarantees that the set that has just been increased has no more elements than the one that has been decreased. That

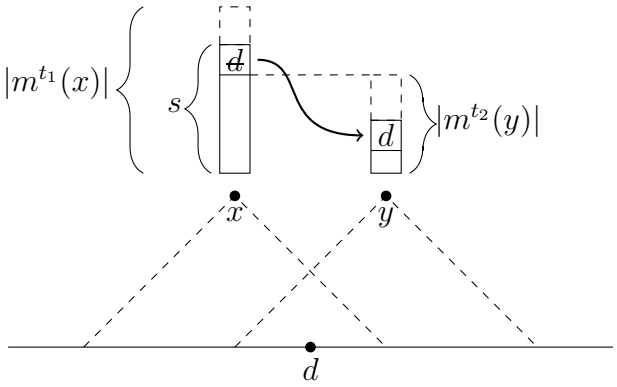

FIGURE 2. The move of element d from $m(x)$ into $m(y)$.

property, and the fact that no vertex that was ready for y had assigned set greater than s, gives $s > |m^{t_2}(y)|$. The claim follows. □

Claim 3.3. If t is the presenting time of u and $N(u) \cap X \neq \emptyset$, then $|m^t(u)| = \alpha$.

Proof. Let $d \in N(u) \cap X$. By the definition of X element d is strongly available for u when u is presented. But d is not chosen in the line 2 of α-BALANCED. It means there were at least α other strongly available elements for u which were added to $m^t(u)$. Thus, indeed $|m^t(u)| = \alpha$. □

Claim 3.4. For any subset $Y \subseteq \mathbf{Y}$ we have

$$(|Q| - |Y|)(\mu - 1) + |M| \leqslant |D - X|\,,$$

where $\mu = \min\{|m(y)| : y \in Y\}, M = \bigcup_{y \in Y} m(y)$ and Q is the set of all vertices $q \in U$ for which $N(q) \cap (M \cup X) \neq \emptyset$.

Proof. The claim is obvious for $\mu = 1$ since $M \cap X = \emptyset$. We assume that $\mu > 1$. Let $i = |Y|\,, s = |Q|$ and let $(q_1, \ldots, q_s)$ be the enumeration of Q for which the sequence of corresponding presenting times $(t_1, \ldots, t_s)$ is strictly decreasing. It means that $q_1, \ldots, q_s$ are in the reverse of the arrival order. For each q_j we recursively define set

$$Z_j := m^{t_j}(q_j) - (Z_1 \cup \ldots \cup Z_{j-1}) \subseteq D - X.$$

Observe that $|Z_1| + \ldots + |Z_s| \leqslant |D - X|$ and to finish the proof it suffices to show that $|Z_j| \geqslant \mu - 1$, for all j $(1 \leqslant j \leqslant s)$, and $Z_j \supseteq m(q_j)$ for $q_j \in Y$. We check two possibilities, either $q_j \in Y$ or $q_j \notin Y$. The first one is straightforward. Set $m^t(q_j)$ can only get smaller, after vertex q_j has been presented, therefore we have $m(q_j) \subseteq m^t(q_j)$ for all $t \geqslant t_j$. In particular, for every $j' \leqslant j - 1$, we have $t_{j'} \geqslant t_j$, hence $m(q_j) \cap Z_{j'} = \emptyset$. That gives $Z_j \supseteq m(q_j)$.

Suppose now, that $q_j \notin Y$ and let $m^{t_j}(q_j) = Z_j \cup R_j$, where $R_j \subseteq Z_1 \cup \ldots \cup Z_{j-1}$. Assume also that $|Z_j| < \alpha$ since otherwise $|Z_j| = \alpha > \mu - 1$. We consider two cases:

Case 1: $R_j = \emptyset$**.** By Claim 3.3 and the definition of $q_j \in Q$, inequality $|Z_j| < \alpha$ implies that $N(q_j) \cap M \neq \emptyset$. It means that some element $d \in M$ was available for q_j at the time when q_j was presented. Element d must belong to $m^{t_j}(u)$ for some $u \in U$ presented earlier (otherwise the algorithm would put it in $m^{t_j}(q_j)$, but since after the game we

have $d \in M$, it would imply that $d \in R_j$). Thus, there are two elements (possibly the same) $u \in Q$ and $y \in Y$ such that $d \in m^{t_j}(u) \cap m(y)$ and obviously y is presented after u. While u is presented before q_j element y may be presented after q_j. Regardless of whether or not u and y are the same by Claim 3.2 we deduce $|m^{t_j}(u)| \geqslant |m(y)| \geqslant \mu$. On the other hand the algorithm in round t_j did not choose element d to be assigned to q_j, which means that at the end of the round the inequality in the line 3 of the algorithm was not satisfied. That means that $|Z_j| = |m^{t_j}(q_j)| \geqslant |m^{t_j}(u)| - 1 \geqslant \mu - 1$.

Case 2: $R_j \neq \emptyset$**.** Let $t > t_j$ be the smallest number (the first moment) for which $R_j \cap m^t(q_j) = \emptyset$ (it is a straightforward consequence of the definition that such t exists). Clearly $|Z_j| \geqslant |m^t(q_j)|$. Consider any $d \in m^{t-1}(q_j) \cap R_j$ and note there is $l < j$ such that $d \in Z_l \subseteq m^{t_l}(q_l)$ with $t - 1 < t_l$. By Claim 3.2 it means that $|m^t(q_j)| \geqslant |m^{t_l}(q_l)|$, thus $|Z_j| \geqslant |Z_l|$. Straightforward induction (with Case 1 as basis) gives $|Z_j| \geqslant \mu - 1$. □

We are ready to prove the proposition. Fix any optimal (maximum) matching in graph G and let $F \subseteq D$ be the set of all elements in D outside the matching. Consider an enumeration $(y_1, \ldots, y_k)$ of $\mathbf{Y}$ such that, for $x_i = |m(y_i)| - 1$, we have $x_1 \geqslant x_2 \geqslant \ldots \geqslant x_k \geqslant 0$. Let $x' = |X - F|$, $f' = |F - X|$ and $x_0 = x' - f'$. Observe that $|F| = f' + |X| - x'$. It implies

$$|D - X| = |D| - |X| = n + |F| - |X| = n - x_0. \tag{4}$$

To show that $(k, (x_0, x_1, \ldots, x_k))$ satisfy (1), fix $1 \leqslant i \leqslant k$ and apply Claim 3.4 for $Y = \{y_1, \ldots, y_i\}$. Then $|M| = x_1 + \ldots + x_i + i$ and $\mu = x_i + 1$. Recall that in the chosen optimal matching each vertex in $D - F$ has a unique match in U. Therefore, $|Q| \geqslant |M - F| + |X - F| = |M - (F - X)| + |X - F| \geqslant |M| - f' + x'$ since $M \cap X = \emptyset$, and then

$$(|M| - i + x_0)(\mu - 1) + |M| \leqslant |D - X| \overset{(4)}{=} n - x_0,$$

which can be rewritten into

$$(x_0 + x_1 + \ldots + x_i) \cdot (1 + x_i) \leqslant n - i.$$

Let Q' be the set of all vertices $q \in U$ for which $N(q) \cap X \neq \emptyset$. Next, we define $s(q)$ as the set of all strong available elements assigned to $m(q)$ in the line 2 of the algorithm. Observe that $s(q_1)$ and $s(q_2)$ are disjoint for distinct $q_1, q_2 \in U$. By Claim 3.3 we get $|s(q)| = \alpha$ for each $q \in Q'$. Thus, $\alpha\,|Q'| \leqslant n - x_0$ since $\bigcup_{q \in Q'} s(q) \subseteq D - X$ and by (4). Also, since each element in $D - F$ has a unique match in U, we have $|Q'| \geqslant |X - F| = x' \geqslant x_0$. Therefore $(1 + \alpha)x_0 \leqslant n$.

To finish the proof recall that $\bigcup_{y \in \mathbf{Y}} m(y) = D - X$. Thus, $x_1 + \ldots + x_k + k = n - x_0$ and consequently the size of the matching constructed by the algorithm equals $k = n - (x_0 + x_1 + \ldots + x_k)$. Also, since k cannot be larger then n we have $x_0 + \ldots + x_k \geqslant 0$. □

Combining Proposition 2.2 and Proposition 3.1 we finally get Theorem 1.1.

4. Competitiveness of α-BALANCED algorithm

Let $\mathrm{bal}(\alpha, n)$ be the worst (minimum) size of matching constructed by α-BALANCED in any α-lazy matching game of size n. Competitive ratio of α-BALANCED is defined as $\mathrm{bal}(\alpha) = \liminf_{n \to \infty} \mathrm{bal}(\alpha, n)/n$. Propositions 2.2 and 3.1 imply that in order to determine $\mathrm{bal}(\alpha, n)$ it is enough to find a pair $(k, (x_0, \ldots, x_k))$ satisfying (1) which maximizes $\sum_{i=0}^{k} x_i$. Moreover, by Proposition 2.1 we can assume that in the maximizing solution we

have $x_0 = \lfloor \frac{n}{1+\alpha} \rfloor$. From now on we consider x_0 in system (1) as fixed together with n and α. Suppose that pair $(k, (x_1, \ldots, x_k))$ satisfies (1). Note that, for $i \geqslant 1$, if $x_i = x_{i+1}$ then the $(i+1)$-th inequality of system (1) implies the i-th inequality. That suggests another representation of the solutions. For a pair (k, x) satisfying (1) with $x = (x_1, \ldots, x_k)$, let $Y(x) = (y_1, \ldots, y_m)$ be such that $m = 1 + x_1$ and $y_j = |\{i > 0 : 1 + x_i = j\}|$. Then, for every i for which $x_{i+1} \neq x_i$, inequality

$$(x_0 + x_1 + \ldots + x_i)(1 + x_i) \leqslant n - i$$

can be rewritten into

$$(x_0 + (m-1) \cdot y_m + \ldots + (t-1) \cdot y_t) \cdot t \leqslant n - (y_m + \ldots + y_t),$$

where $t = x_i + 1$. By the above discussion sequence $(y_1, \ldots, y_m)$ belongs to the image of Y whenever it satisfies the following system of inequalities

$$\Psi_{n,m}(x_0): \qquad t \cdot x_0 + \sum_{i=t}^{m} (1 + (i-1)t)\mathbf{y_i} \leqslant n \quad \text{for } t = 1, \ldots, m \tag{5}$$

Moreover, since $y_i \geqslant 0$ then m-th inequality in (5) implies

$$n - mx_0 \geqslant 0. \tag{6}$$

On the other hand, having any $m > 0$ satisfying (6) and any solution $(y_1, \ldots, y_m)$ of (5) one can easily find (using the definition of y_j's) a solution (k, x) of (1) such that $x_0 + x_1 + \ldots + x_k = x_0 + \sum_{i=1}^{m} (i-1)y_i$.

The above considerations are summarized in the following

Proposition 4.1. *The minimal size of the matching constructed by α-BALANCED in all α-lazy matching games of size n is equal to*

$$\mathrm{bal}(\alpha, n) = n - x_0 - \sup\left\{ \sum_{i=1}^{m} (i-1)y_i \right\},$$

where $x_0 = \lfloor n/(1+\alpha) \rfloor$ and supreme is taken over all integers $m > 0$ such that $n - mx_0 \geqslant 0$ and over all vectors $(y_1, \ldots, y_m)$ of nonnegative integers that satisfy $\Psi_{n,m}(x_0)$.

4.1. **LP formulation.** In order to maximize $\sum_{i=1}^{m} (i-1)y_i$ we consider the following linear program. We present the primal and the dual formulation.

- Primal:

maximize	$\sum_{j=1}^{m} (j-1)y_j$	
subject to	$\sum_{j=i}^{m} (1 + (j-1)i)y_j \leqslant n - ix_0$,	for $i = 1, \ldots, m$
and	$y_j \geqslant 0$,	for $j = 1, \ldots, m$

- Dual:

minimize	$\sum_{i=1}^{m} (n - ix_0)z_i$	
subject to	$\sum_{i=1}^{j} (1 + (j-1)i)z_i \geqslant j - 1$,	for $j = 1, \ldots, m$
and	$z_i \geqslant 0$,	for $i = 1, \ldots, m$

Let $P_i(y) = \sum_{j=i}^{m} (1 + (j-1)i)y_j$ and $D_j(z) = \sum_{i=1}^{j} (1 + (j-1)i)z_i$. Observe that both systems $(P_i(y) = n - ix_0)_{i=1,\ldots,m}$ and $(D_j(z) = j - 1)_{j=1,\ldots,m}$ are quadratic and have unique solutions. By the Complementary Slackness Theorem these solutions are optimal if both are nonnegative vectors.

To solve the first system observe that $P_{i+2}(y) + P_i(y) - 2P_{i+1}(y) = (1 - i + i^2)y_i - (1 + i)^2 y_{i+1}$. This, together with initial values for y_m and y_{m-1} gives

$$y_i = y_{i+1}\frac{(i+1)^2}{1-i+i^2} \quad \text{for } i = 1, \ldots, m-2, \tag{7}$$

$$y_m = \frac{n - m \cdot x_0}{1 + m(m-1)}, \qquad y_{m-1} = \frac{x_0 + (m-1)y_m}{1 + (m-1)(m-2)}. \tag{8}$$

By (6) $y_m \geqslant 0$ and thus all y_i are nonnegative.

For the second system observe that $D_{j+1}(z) + D_{j-1}(z) - 2D_j(z) = (1 + j + j^2)z_{j+1} - (1 - j)^2 z_j$. Again, with the initial values for z_1 and z_2 we get

$$z_{j+1} = z_j \frac{(j-1)^2}{1+j+j^2} \quad \text{for } j = 2, \ldots, m-1,$$
$$z_1 = 0, \qquad z_2 = 1/3.$$

Therefore $z_i \geqslant 0$ and indeed the above vectors y and z are optimal solutions for the primal and the dual LP.

To calculate the target function $\sum_{j=1}^{m}(j-1)y_j$ consider new variable y_0 defined by the additional condition: $P_0(y) = n$. Equation (7) still works for y_0. Also, after combining $P_0(y)$ with $P_1(y)$ we get $\sum_{j=1}^{m}(j-1)y_j = y_0 - x_0$. Finally, using (7) and (8) and after some rearrangement we find

$$x_0 + \sum_{j=1}^{m}(j-1)y_j = y_0 = \frac{(m-1)n + x_0}{m} \prod_{i=1}^{m-1} \frac{i+i^2}{1+i+i^2}. \tag{9}$$

4.2. **Lower bound.** The solution for LP relaxation of (5) may not be integer therefore with the solution (9) we have only the following bound

$$\mathrm{bal}(\alpha, n) \geqslant n - \sup\{y_0 : n - mx_0 \geqslant 0\}.$$

Let $F(z, m) = \frac{(m-1)+z}{m}\prod_{i=1}^{m-1}\frac{i+i^2}{1+i+i^2}$. Then, for fixed m, we get $y_0 = n \cdot F(x_0/n, m)$. Observe that function $F(z, m)$ is increasing with m, for $m \leqslant 1/z$ and $F(0, m)$ increases indefinitely with m. Therefore for $x_0 = \lfloor n/(1+\alpha) \rfloor > 0$, we have

$$\begin{aligned} \mathrm{bal}(\alpha, n)/n &\geqslant 1 - \max_{m \leqslant n/x_0} F(x_0/n, m) \\ &= 1 - F(x_0/n, \lfloor n/x_0 \rfloor) \overset{n\to\infty}{\longrightarrow} 1 - F(1/(1+\alpha), 1+\alpha), \end{aligned} \tag{10}$$

and for $x_0 = 0$ (it happens when $\alpha \geqslant n$) the bound is

$$\mathrm{bal}(\alpha, n)/n \geqslant 1 - \lim_{m\to\infty} F(0, m) = 1 - \prod_{i=1}^{\infty}\frac{i+i^2}{1+i+i^2} = 1 - \frac{\pi}{\cosh\frac{\sqrt{3}}{2}\pi}. \tag{11}$$

4.3. **Upper bound.** Consider $m = \min(\lfloor n/x_0 \rfloor, \ln n)$ assuming n/x_0 is greater then $\ln n$ when $x_0 = 0$, and let $(y_1, \ldots, y_m)$ be the optimal rational solution of $\Psi_{n,m}(x_0)$. Let $v = (v_1, \ldots, v_m)$ be such that $v_i = \lfloor y_i \rfloor$. Vector v contains only nonnegative,

integer entries and $v \leqslant y$. The shape of system $\Psi_{n,m}(x_0)$ guarantees that v also satisfies $\Psi_{n,m}(x_0)$. Finally we have

$$x_0 + \sum_{i=1}^{m}(i-1)v_i > x_0 + \sum_{i=1}^{m}(i-1)y_i - \sum_{i=1}^{m}(i-1).$$

By the definition of m we know that $m = o(n)$. Therefore by Proposition 2.2 we get

$$\mathrm{bal}(\alpha, n) < n \cdot (1 - F(x_0/n, m)) + m(m-1)/2 = n \cdot (1 - F(x_0/n, m)) + o(n)$$

Hence, for $x_0 = \lfloor n/(1+\alpha) \rfloor > 0$ the bound is

$$\mathrm{bal}(\alpha, n)/n < 1 - F(x_0/n, \lfloor n/x_0 \rfloor) + o(1) \overset{n\to\infty}{\longrightarrow} 1 - F(1/(1+\alpha), 1+\alpha), \tag{12}$$

and for $x_0 = 0$ we have

$$\mathrm{bal}(\alpha, n)/n < 1 - F(0, \ln n) + o(1) \overset{n\to\infty}{\longrightarrow} 1 - \prod_{i=1}^{\infty} \frac{i+i^2}{1+i+i^2} = 1 - \frac{\pi}{\cosh \frac{\sqrt{3}}{2}\pi}. \tag{13}$$

4.4. **Competitive ratio.** When α is finite, by (10) and (12) the ratio of α-BALANCED is equal to $\mathrm{bal}(\alpha) = 1 - F(\frac{1}{\alpha+1}, \alpha+1)$. For the case $\alpha \to \infty$ note that $F(1/(1+\alpha), 1+\alpha) \to \pi/\cosh \frac{\sqrt{3}}{2}\pi$. Thus, by all (10-13) we get $\mathrm{bal}(\alpha) = 1 - \pi/\cosh \frac{\sqrt{3}}{2}\pi$. This finally proves Theorem 1.2.

5. Conclusion and remarks

In the classical on-line matching problem, randomized approach has a big advantage over the deterministic one. This paper shows that the lazy method moves forward deterministic bounds. It is interesting to know what can be achieved with randomization of the lazy technique.

Problem 1. What is the competitive ratio of the randomized version of the lazy matching problem?

The authors of [12] consider a variant (called b-matching) where each server can realize up to b tasks. Competitive ratio of their optimal algorithm approaches $1 - \frac{1}{e}$ with $b \to \infty$, which is a barrier for any randomized matching algorithm (see [14]). We expect that the lazy method is capable of breaking $1 - \frac{1}{e}$ limit in case of b-matching.

Problem 2. What is the competitive ratio of the lazy b-matching problem?

References

[1] Yossi Azar and Yoel Chaiutin. Optimal node routing. In Bruno Durand and Wolfgang Thomas, editors, *STACS 2006*, volume 3884 of *Lecture Notes in Computer Science*, pages 596–607. Springer Berlin Heidelberg, 2006.

[2] Yossi Azar and Yossi Richter. Management of multi-queue switches in qos networks. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, STOC '03, pages 82–89, New York, NY, USA, 2003. ACM.

[3] Bahman Bahmani and Michael Kapralov. Improved bounds for online stochastic matching. In *Proceedings of the 18th annual European conference on Algorithms: Part I*, ESA'10, pages 170–181, Berlin, Heidelberg, 2010. Springer-Verlag.

[4] Benjamin Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *SIGACT News*, 39(1):80–87, March 2008.

[5] Ashish Chiplunkar, Sumedh Tirodkar, and Sundar Vishwanathan. On randomized algorithms for matching in the online preemptive model. In Nikhil Bansal and Irene Finocchi, editors, *Algorithms ESA 2015*, volume 9294 of *Lecture Notes in Computer Science*, pages 325–336. Springer Berlin Heidelberg, 2015.

[6] Nikhil R. Devenur and Thomas P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*, EC '09, pages 71–78, New York, NY, USA, 2009. ACM.

[7] Leah Epstein, Asaf Levin, Danny Segev, and Oren Weimann. Improved bounds for online preemptive matching. In *30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013, February 27 - March 2, 2013, Kiel, Germany*, pages 389–399, 2013.

[8] Jon Feldman, Nitish Korula, Vahab Mirrokni, S. Muthukrishnan, and Martin Pl. Online ad assignment with free disposal. In Stefano Leonardi, editor, *Internet and Network Economics*, volume 5929 of *Lecture Notes in Computer Science*, pages 374–385. Springer Berlin Heidelberg, 2009.

[9] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating 1-1/e. In *Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '09, pages 117–126, Washington, DC, USA, 2009. IEEE Computer Society.

[10] Stefan Felsner. On-line chain partitions of orders. *Theoret. Comput. Sci.*, 175(2):283–292, 1997.

[11] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, pages 982–991, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.

[12] Bala Kalyanasundaram and Kirk R. Pruhs. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science*, 233(1-2):319 – 325, 2000.

[13] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, STOC '11, pages 587–596, New York, NY, USA, 2011. ACM.

[14] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, STOC '90, pages 352–358, New York, NY, USA, 1990. ACM.

[15] Kamil Kloch. Online dimension of partially ordered sets. *Reports on Mathematical Logic*, 42:101–116, 2007.

[16] Vahideh H. Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: online actions based on offline statistics. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 1285–1294. SIAM, 2011.

[17] Aranyak Mehta and Debmalya Panigrahi. Online matching with stochastic rewards. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, FOCS '12, pages 728–737, Washington, DC, USA, 2012. IEEE Computer Society.

[18] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5):no. 5, October 2007.